

Development of A Platform for Large-scale Reservoir Simulations on Parallel computers

Hui Liu*, Kun Wang, Bo Yang, and Zhangxin Chen
 Department of Chemical and Petroleum Engineering
 University of Calgary
 Calgary, Alberta, Canada, T2N 1N4

Abstract

This paper presents our work on designing a platform for large-scale reservoir simulations. Detailed components, such as grid and linear solver, and data structures are introduced, which can serve as a guide to parallel reservoir simulations and other parallel applications. The main objective of platform is to support implementation of various parallel reservoir simulators on distributed-memory parallel systems, where MPI (Message Passing Interface) is employed for communications among computation nodes. It provides structured grid due to its simplicity and cell-centered data is applied for each cell. The platform has a distributed matrix and vector module and a map module. The matrix and vector module is the base of our parallel linear systems. The map connects grid and linear system modules, which defines various mappings between grid and linear systems. Commonly-used Krylov subspace linear solvers are implemented, including the restarted GMRES method and the BiCGSTAB method. It also has an interface to a parallel algebraic multigrid solver, BoomerAMG from HYPRE. Parallel general-purpose preconditioners and special preconditioners for reservoir simulations are also developed. Various data structures are designed, such as grid, cell, data, linear solver and preconditioner, and some key default parameters are presented in this paper. The numerical experiments show that our platform has excellent scalability and it can simulate giant reservoir models with hundreds of millions of grid cells using thousands of CPU cores.

Keywords: platform, reservoir simulation, parallel computing, algorithm, data structure

1 Introduction

Nowadays, various operation processes have been developed to enhance oil recovery by the oil and gas industry. Their numerical simulations are becoming more and more complicated. In the meantime, geological models from reservoirs are more and more complex, and they are also heterogenous. Models with millions of grid cells are usually employed to obtain high resolution results. Numerical simulations may take days or even longer to complete one run using regular workstations. The long simulation time could be a problem to reservoir engineers, since dozens of simulations may be required to find optimal operations. Fast computational methods and reservoir simulators should be investigated.

*Authors to whom correspondence may be addressed. Email addresses: hui.jw.liu@gmail.com

Reservoir simulations have been studied for decades and various models and methods have been developed by researchers, including black oil model, compositional model, thermal model and related topics. Kaarstad et al. [6] studied oil-water model and they implemented a reservoir simulator that could solve problems with up to one million grid cells. Rutledge et al. [4] developed a compositional simulator for massive SIMD computers, which employed the IMPES (implicit pressure-explicit saturation) method. Killough et al. [3] implemented a compositional simulator for distributed-memory parallel systems. Killough et al. also used the locally refined grids in their parallel simulator to improve accuracy [7]. Dogru and his group [8, 9] developed a parallel simulator, which was capable of simulating reservoir models with one billion grid cells. Zhang et al. developed a platform for adaptive finite element and adaptive finite volume methods, which has been applied to CFD, Maxwell equation, material, electronic structures, biology and reservoir simulations [10, 11, 39], and a black oil simulator using discontinuous Galerkin method has been reported [39]. For many reservoir simulations, especially black oil simulation, most of the simulation time is spent on the solution of linear systems and it is well-known that the key of accelerating linear solvers is to develop efficient preconditioners. Many preconditioner methods have been applied to reservoir simulations, including point-wise and block-wise incomplete factorization (ILU) methods for general linear systems [13], domain decomposition methods [22], constrained pressure residual (CPR) methods for the black oil model, compositional model and extended black oil models [14, 15], multi-stage methods [17], multiple level preconditioners [38] and fast auxiliary space preconditioners (FASP) [18].

This paper presents our work on developing a parallel platform for large-scale reservoir simulations on parallel systems and designing various data structures. The platform is implemented using C and MPI (Message Passing Interface). MPI is a standardized message-passing system designed to work on a wide varieties of parallel system and it is employed to handle communications among computation nodes. The platform provides structured grid, cell-centered data, linear solvers, preconditioners, distributed matrices and vectors, visualization, parallel input and output through MPI-IO, key words parsing and well modeling modules. Finite difference methods and finite volume methods are supported. The load balancing module is crucial for parallel computing [34, 36, 24]. The load balancing module is completed by ParMETIS [36] and the Hilbert space-filling curve (HSFC) method [34]. The ParMETIS is a graph partitioning package using topological information of a grid, and the HSFC partitioning method is an in-house partitioning method, which serves as the default partitioner. Commonly used Krylov subspace solvers and algebraic multigrid (AMG) solvers are implemented, including the restarted GMRES solver, BiCGSTAB solver [12], and classic AMG solvers [33]. General preconditioners, including ILU(k), ILUT, domain decomposition [22] and AMG [33], and special preconditioners, including CPR-like preconditioners, are implemented. Detailed designs and key parameters are presented. Numerical experiments show that our platform is capable of calculating problems with hundreds of millions of grid cells and it has excellent scalability on distributed-memory parallel computers.

2 Data Types

Data structures in this paper use C language style. Figure 1 presents our basic data types for integer, float-point, boolean and character types and MPI-related data types used by MPI. These data types are defined automatically by `configure`, which is generated by `autoconf`, `m4` and `BASH` scripts.

The default type for floating-point number is `double`. If long double is enabled, its type is `long double`. Macro `USE_LONG_DOUBLE` is defined to define derived types, such as `PRSI_MPI_FLOAT` and `COORD`. The default integer type is `long int`. Two macros, `USE_LONG_LONG` and `USE_LONG`, are

```

/* float */
#if USE_LONG_DOUBLE
typedef long double      FLOAT;
#else
typedef double           FLOAT;
#endif

/* integer */
#if USE_LONG_LONG
typedef signed long long int  INT;
#elif USE_LONG
typedef signed long int      INT;
#else
typedef signed int           INT;
#endif

typedef char              CHAR;
typedef FLOAT              COORD[3];
#undef TRUE
#define TRUE                (1)
#undef FALSE
#define FALSE                (0)
typedef int                BOOLEAN;

/* MPI type */
#if USE_LONG_DOUBLE
#define PRSI_MPI_FLOAT      MPI_LONG_DOUBLE
#else
#define PRSI_MPI_FLOAT      MPI_DOUBLE
#endif

#if USE_LONG_LONG
#define PRSI_MPI_INT        MPI_LONG_LONG_INT
#elif USE_LONG
#define PRSI_MPI_INT        MPI_LONG
#else
#define PRSI_MPI_INT        MPI_INT
#endif

#define PRSI_MPI_CHAR       MPI_CHAR
#define PRSI_MPI_BOOLEAN    MPI_INT

/* CSR format */
typedef struct mat_csr_t_
{
    INT    num_rows;
    INT    num_cols;
    INT    num_nonzeros;
    INT    *Ap;
    INT    *Aj;
    FLOAT  *Ax;
} mat_csr_t;

```

Figure 1: Basic data types

defined to control integer types and derived data types, such as MPI types. `BOOLEAN` is for boolean type and `CHAR` is for string. `mat_csr_t` is basic data type for CSR matrix.

3 Grid

A traditional reservoir can be described as $\Omega = [x_1, x_2] \times [y_1, y_2] \times [z_1, z_2]$. If the domain Ω are divided into n_x , n_y and n_z intervals in the x , y and z directions, then the grid has $N_g = n_x \times n_y \times n_z$ cells. Each cell is a hexahedron. An interior cell has six neighbors and each boundary cell may have three, four or five neighbors, depending on its location. A structured grid can be uniform or non-uniform. The structured grids support finite difference methods, finite volume methods and finite element methods, and they have been widely used by commercial reservoir simulators.

Each cell has a unique global index. Its default index is calculated as

$$C_{(i,j,k)} = n_x * n_y * k + n_x * j + i, \quad (1)$$

which is numbered from the bottom layer of a reservoir to the top layer of the reservoir. Here i , j and k are the integer coordinates of the cell in the x , y and z directions, respectively. Another numbering style used by most reservoir simulators is

$$C_{(i,j,k)} = n_x * n_y * (n_z - k) + n_x * j + i, \quad (2)$$

which is numbered from top layer to bottom layer.

```
typedef struct CELL_
{
#if !USE_LESS_MEMORY
    COORD    ctrd;           /* centroid coordinate */
    FLOAT    area[6];       /* area of each face */
    FLOAT    vol;           /* volume */
#endif

    void      *nb[6];        /* neighbours */
    INT        vert[8];      /* local index of vertices */
    INT        index;        /* local index */
    INT        idx[3];       /* index in x, y and z direction */
    INT        regn;         /* region mark */

    USHORT     bdry_type[6]; /* boundary type */
    USHORT     type;         /* cell type */
} CELL;
```

Figure 2: Data structure of CELL

The data structure of a cell, which is defined as `CELL`, is shown by Figure 2. This data structure defines cell-related information, such as centroid coordinate (`ctrd`), area of all faces (`area`) and volume (`vol`), six neighbors (`nb`), local index of each vertex (`vert`), local index on an MPI process (`index`), global index in three directions (`idx`), region mark (`regn`), boundary type of each face (`bdry_type`) and cell type (`type`). Some members are optional. For example, if we would like to use less memory, we can remove `ctrd`, `area` and `vol` by setting `USE_LESS_MEMORY` to some positive integer, such as 1. In reservoir simulations, each cell represents a portion of reservoir and they have similar properties, such as porosity, pressure, temperature, water saturation and oil saturation.

```
typedef struct RNEIGH_
{
    INT      gidx;
    INT      lidx;
    int      rank;
} RNEIGH;
```

Figure 3: Data structure of remote neighbor

When a neighbor of a cell is in another MPI process, communication is required when accessing neighbor information. The data structure **RNEIGH** stores remote cell information, such as its global cell index (**gidx**), its local index (**lidx**), and its MPI rank (**rank**). The member **rank** defines message sender and receiver during communication.

Data structure for structured grids, **GRID**, is presented in Figure 4, which stores the coordinates of each vertex (**vert**), vertex indices (**L2Gmap_vert**), cell indices (**L2Gmap_cell**), distribution of cells (**num_cells**) in each MPI process, a mapping between the global index of a cell and its local index, well data and MPI info. In each MPI process, a portion of a grid is stored, and only vertices and faces belong to these cells are stored. Space cost of a grid in each MPI process is proportional to grid size (number of cells).

3.1 Grid Partitioning

Let \mathbb{G} be the structured grid, which is distributed in N_p MPI tasks,

$$\mathbb{G} = \{C_1, C_2, \dots, C_{N_g}\}, \quad (3)$$

where C_i is the i -th cell of \mathbb{G} . Each task owns a subset of \mathbb{G} , \mathbb{G}_i , which satisfies the following conditions:

$$\begin{cases} \mathbb{G}_i \neq \emptyset \ (i = 1, \dots, N_p) \\ \mathbb{G}_i \cap \mathbb{G}_j = \emptyset \ (i \neq j) \\ \cup \mathbb{G}_i = \mathbb{G} \ (i = 1, \dots, N_p). \end{cases} \quad (4)$$

A cell belongs to some sub-grid, and its neighboring cells may belong to different sub-grids.

Each cell has similar calculations. Therefore, it is reasonable to assume that each cell has the same amount of calculations. The workload of each MPI task can be modeled by the size of its sub-grid, $|\mathbb{G}_i|$, or simply the number of grid cells in the sub-grid. When discretizing reservoir models, information from neighboring cells is always required. For any sub-grid, its communication volume is determined by its remote neighboring cells, which can be modeled by dual graph: we see a cell as a vertex of a graph, and if two cells are neighbors, there exists an edge between these two cells.

The goal of grid partitioning is that each MPI task has equal workload (number of cells) and the communications are minimized. Graph methods are ideal tools, such as spectral methods [41, 42], multilevel methods [43, 44, 45] and diffusive methods. Several graph partitioning packages have been implemented and available publicly, such as ParMETIS, which is widely applied in parallel computing. It is also used for matrix reordering. The geometry information based methods are also efficient, including recursive coordinate bisection method, recursive inertial bisection method and the space-filling curve methods [34].

```

typedef struct GRID_
{
    COORD          *vert;           /* coordinates of each vertex */
    CELL           *cell;
    INT             *num_cells;      /* number of cells in each process */
    RNEIGH          *rnghr;         /* remote neighbours */
    USHORT          *type_vert;     /* vert types */
    INT             *L2Gmap_vert;    /* Local to global map of vertices */
    INT             *L2Gmap_cell;   /* Local to global map of cell indices */

    FLOAT           lif;            /* Load imbalance factor */
    INT             nregns;         /* number of region marks */
    INT             nverts;         /* number of vertices in the subgrid */
    INT             nfaces;         /* nuber of faces */
    INT             ncells;         /* number of cell indices in the subgrid */
    INT             nrngbr;         /* number of remote neighbours */

    INT             nfaces_remote;  /* equals to number of remote neighbours */
    INT             nverts_global;  /* number of vertices in the global grid */
    INT             nfaces_global;  /* number of vertices in the global grid */
    INT             ncells_global;  /* number of cells in the global grid */

    FLOAT           bbox[3][2];     /* bounding box */
    INT             ncx, ncy, ncx;  /* grid size in x, y, z directions */
    FLOAT           *vx, *vy, *vz; /* partition of x, y, and z directions */
    BOOLEAN         uniform;        /* uniform in each direction or not */

    /* Well, (nprocs - 1)-th process owns all wells */
    WELL            **well;
    WELL_CINFO      *well_cinfo;
    CELL            **perf_cell;    /* pointer to cell which has perferation */
    INT             nperfs;         /* number of perferations */
    INT             nperfs_global;
    INT             nwells_global;  /* number of wells */
    BOOLEAN         well_assembled;
    BOOLEAN         destroy_well;   /* if grid or simulator destroy wells */

    MPI_Comm        comm;
    int             rank;
    int             nprocs;
} GRID;

```

Figure 4: Data structure of GRID

3.2 Space-filling Partitioning Method

This section introduces space-filling curve, Hilbert order, algorithm for generating Hilbert order and space-filling curve partitioning method.

3.2.1 Space-filling Curves

Space-filling curves are those curves that fill an entire n -dimensional unit hypercube, which were proposed by Peano in 1890 and popularized by Hilbert later.

Many space-filling curves have been discovered. Figures 5 and 6 show levels 1, 2, 3 and 6 Hilbert




Figure 5: Hilbert space-filling curves, two dimensions, levels 1, 2 and 3




Figure 6: Hilbert space-filling curve, level 6

space-filling curves in a two-dimensional unit square. We can see that a curve is denser if a level is higher. Figure 7 shows levels 1, 2 and 3 Sierpiński space-filling curves. Figure 8 shows a level 2 Morton space-filling curve. From these curves, we can observe that the Hilbert space-filling curves and the Sierpiński space-filling curves have good locality, and the Morton space-filling curves have jumps, whose locality is poor.

3.2.2 Hilbert Orders

Each curve has a starting point and an ending point. Along this curve, a map is introduced between a one-dimensional domain and a multi-dimensional domain. Figure 9 and 10 show a level 2 Hilbert




Figure 7: Sierpiński space-filling curves, levels 1, 2 and 3




Figure 8: Morton space-filling curve, level 2

space-filling curve and a level 2 Morton space-fill curve (Hilbert order), respectively. Both of them have 16 vertices whose indexes start from 0 to 15. We can see that these curves define orders, which map a two-dimensional space to a one-dimensional space. Higher dimensional space-filling curves are defined similarly.

A Hilbert space-filling curve is one famous space-filling curve, which is also called a Hilbert curve [60]. This kind of curve has many important characteristics, such as locality, clustering and self-similarity. A Hilbert curve (order) has been applied in many areas, including image storing, database indexing, data compression and dynamic load balancing. For parallel computing, the Hilbert order method is one of the most important geometry-based partitioning methods.




Figure 9: Hilbert order, level 2




Figure 10: Morton order (Z-order), level 2

An n ($n \geq 2$)-dimensional Hilbert curve introduces a one-to-one mapping between an n -dimensional space and a one-dimensional space. The mapping from the n -dimensional space to the one-dimensional space is called encoding, while the inverse mapping is called decoding, which

maps an integer to a coordinate in the n -dimensional space. Algorithms for computing Hilbert curves/orders in two- and three-dimensional spaces have been proposed in the literature, which can be classified into recursive algorithms [56, 57, 58, 59] and iterative algorithms [47, 46, 52, 53, 55, 50]. Iterative algorithms, especially the table-driven algorithms [47, 46], are usually much faster than recursive algorithms. In general, the complexities of these algorithms are $O(m)$, where m is the level of a Hilbert curve. For a two-dimensional space, Chen et al. [50] proposed an algorithm of $O(r)$ complexity, where r is defined as $r = \log_2(\max(x, y)) + 1$, $r \leq m$ and is independent of the level m . This algorithm is faster when m is much larger than r . The same idea was also applied to a three-dimensional space [34]. For higher dimensional spaces, Kamata et al. presented a representative n -dimensional Hilbert mapping algorithm [51] and Li et al. introduced algorithms for analyzing the properties of n -dimensional Hilbert curves [49]. Liu et al. introduced high-order encoding and decoding algorithms [48].

3.2.3 Calculation of Arbitrary Dimensional Hilbert Orders

For the sake of completeness, calculations of Hilbert orders are introduced in this paper [48]. Let m be the level of the Hilbert curves (orders) and D_m be the coordinate set of the m th level Hilbert curve, where $D_m = \{(x_n, \dots, x_2, x_1) | 0 \leq x_i < 2^m, 1 \leq i \leq n\}$. $(x_n, \dots, x_1) \in D_m$ is a coordinate of the Hilbert curve, and x_i ($1 \leq i \leq n$) is called the i th component of the coordinate. A logical operation, **exclusive or operation (xor)**, \wedge for two coordinates is defined as

$$(x_n, \dots, x_2, x_1) \wedge (y_n, \dots, y_2, y_1) = (x_n \wedge y_n, \dots, x_2 \wedge y_2, x_1 \wedge y_1). \quad (5)$$

Here we use $(a_1 a_2 \dots a_k)_d$ to represent a number, where $0 \leq a_i < d$ ($1 \leq i \leq k$), and k can be any positive integer. This number is a binary number if we set $d = 2$ and a decimal number if we set $d = 10$. The **and** operator is denoted by $\&$. Let us define

$$p_n^i(a_1, a_2, \dots, a_n) = \left(\sum_{j=1}^i a_j \right) \bmod 2 = \left(\sum_{j=1}^i a_j \right) \% 2, \quad (6)$$

where a_i equals 0 or 1 ($1 \leq i \leq n$). The $p_n^i(a_1, a_2, \dots, a_n)$ equals 1 or 0. With the help of p_i , function f_n is defined as

$$f_n(a_1, a_2, \dots, a_n) = (b_1 b_2 \dots b_n)_2 = j, b_1 = a_1, b_i = \begin{cases} a_i & \text{if } p_n^{i-1} = 0 \\ 1 - a_i & \text{if } p_n^{i-1} = 1 \end{cases}, \quad (7)$$

where a_i equals 0 or 1 and j is a decimal number. This function f_n maps a vector (a_1, a_2, \dots, a_n) to a decimal number j . Its inverse function b_n is defined as

$$b_n(j) = b_n((a_1 a_2 \dots a_n)_2) = (b_1, b_2, \dots, b_n), b_1 = a_1, b_i = \begin{cases} a_i & \text{if } a_{i-1} = 0 \\ 1 - a_i & \text{if } a_{i-1} = 1 \end{cases}, \quad (8)$$

where j ($0 \leq j < 2^n$) is a decimal number and $j = (a_1 a_2 \dots a_n)_2$. This function maps a decimal number (scalar) to a vector.

For any point x , it can be mapped to an integer vector such that $(x_n, \dots, x_1) \in D_m$. Each component x_i ($1 \leq i \leq n$) is written as $x_i = (x_i^m x_i^{m-1} \dots x_i^1)_2$. The calculated Hilbert order is stored as $(r_m r_{m-1} \dots r_1)_{2^n}$, which will be mapped to $[0, 1]$ in the end. The mapping is described in Algorithm 1, which is an iterative method. Here we assume that $G_n^{i,0}$ and $G_n^{i,1}$ are known, which can be calculated by the algorithms introduced in [49]. This algorithm is for arbitrary dimensional cases. For a specific dimension, such as two dimensions, special optimization techniques can be applied to accelerate calculations [48].

Algorithm 1 Hilbert order algorithm

- (1) For $x \in [0, 1]^n$, map x to domain D_m by multiplying 2^m . Let $k = m$.
 - (2) If $k = 0$, terminate the procedure. Or we have $r_k = f_n(x_n^k, x_{n-1}^k, \dots, x_1^k)$.
 - (3) For each integer i ($1 \leq i \leq n$), if x_i^k equals 1, then $x_i = x_i - 2^{k-1}$.
 - (4) For each integer i ($1 \leq i \leq n$), if the i th component of $G_n^{r_k, 1}$ is 1, then $x_i = 2^{k-1} - 1 - x_i$.
 - (5) If $G_n^{r_k, 0}$ has two components equal to 1 in the i th and j th positions, then swap x_i and x_j .
 - (6) $k = k - 1$, go to (1).
 - (7) $h = \frac{(r_m r_{m-1} \dots r_1) 2^n}{2^{nm}}$ ($h \in [0, 1]$).
-

3.2.4 Table-driven Algorithms for Low Dimensional Hilbert Orders

Table-driven algorithms were introduced in [47, 46]. The basic idea is to store additional information other than to calculate. More memory is required, but computations are faster.

```
typedef double FLOAT;
typedef int INT;
typedef int BOOLEAN;

typedef struct HSFC_ENTRY_ {
    FLOAT coord[3];
    FLOAT hsfc;
} HSFC_ENTRY;

static INT hsfc_maxlevel = 30;
```

Figure 11: Data types for Hilbert order

```
static unsigned const int idata3d[] = {
    0, 7, 3, 4, 1, 6, 2, 5, 0, 1, 3, 2, 7, 6, 4, 5,
    0, 3, 7, 4, 1, 2, 6, 5, 2, 3, 5, 4, 1, 0, 6, 7,
    4, 5, 3, 2, 7, 6, 0, 1, 4, 7, 3, 0, 5, 6, 2, 1,
    6, 7, 5, 4, 1, 0, 2, 3, 0, 1, 7, 6, 3, 2, 4, 5,
    2, 1, 5, 6, 3, 0, 4, 7, 6, 1, 5, 2, 7, 0, 4, 3,
    0, 7, 1, 6, 3, 4, 2, 5, 2, 1, 3, 0, 5, 6, 4, 7,
    4, 7, 5, 6, 3, 0, 2, 1, 4, 5, 7, 6, 3, 2, 0, 1,
    6, 1, 7, 0, 5, 2, 4, 3, 0, 3, 1, 2, 7, 4, 6, 5,
    2, 3, 1, 0, 5, 4, 6, 7, 6, 7, 1, 0, 5, 4, 2, 3,
    2, 5, 1, 6, 3, 4, 0, 7, 4, 3, 7, 0, 5, 2, 6, 1,
    4, 3, 5, 2, 7, 0, 6, 1, 6, 5, 1, 2, 7, 4, 0, 3,
    2, 5, 3, 4, 1, 6, 0, 7, 6, 5, 7, 4, 1, 2, 0, 3
};
```

Figure 12: Ordering table for Hilbert order

In [46], the authors introduced the Gray coding and Morton ordering, which are easy to compute. Other orderings, such as the Hilbert ordering, can be generated through pre-defined appropriate mappings. An ordering table and an orientation table are required to map between Hilbert and Morton orders. Figure 11 shows basic data structures and data types, from which integers, floating-point numbers and Hilbert orders are defined. Figures 12 and 13 are the ordering table and orientation table between the Morton and Hilbert orders. Figure 14 shows internal state conversion

```
static unsigned const int ystate3d[] = {
    1,  6,  3,  4,  2,  5,  0,  0,  0,  7,  8,  1,  9,  4,  5,  1,
    15, 22, 23, 20, 0,  2,  19, 2,  3,  23, 3,  15, 6,  20, 16, 22,
    11, 4,  12, 4,  20, 1,  22, 13, 22, 12, 20, 11, 5,  0,  5,  19,
    17, 0,  6,  21, 3,  9,  6,  2,  10, 1,  14, 13, 11, 7,  12, 7,
    8,  9,  8,  18, 14, 12, 10, 11, 21, 8,  9,  9,  1,  6,  17, 7,
    7,  17, 15, 12, 16, 13, 10, 10, 11, 14, 9,  5,  11, 22, 0,  8,
    18, 5,  12, 10, 19, 8,  12, 20, 8,  13, 19, 7,  5,  13, 18, 4,
    23, 11, 7,  17, 14, 14, 6,  1,  2,  18, 10, 15, 21, 19, 20, 15,
    16, 21, 17, 19, 16, 2,  3,  18, 6,  10, 16, 14, 17, 23, 17, 15,
    18, 18, 21, 8,  17, 7,  13, 16, 3,  4,  13, 16, 19, 19, 2,  5,
    16, 13, 20, 20, 4,  3,  15, 12, 9,  21, 18, 21, 15, 14, 23, 10,
    22, 22, 6,  1,  23, 11, 4,  3,  14, 23, 2,  9,  22, 23, 21, 0
};
```

Figure 13: Orientation table for Hilbert order

```
static unsigned const int *d[] = {
    idata3d,      idata3d + 8,   idata3d + 16,   idata3d + 24,
    idata3d + 32, idata3d + 40,  idata3d + 48,   idata3d + 56,
    idata3d + 64, idata3d + 72,  idata3d + 80,   idata3d + 88,
    idata3d + 96, idata3d + 104, idata3d + 112,  idata3d + 120,
    idata3d + 128, idata3d + 136, idata3d + 144,  idata3d + 152,
    idata3d + 160, idata3d + 168, idata3d + 176,  idata3d + 184
};

static unsigned const int *s[] = {
    ystate3d,      ystate3d + 8,   ystate3d + 16,   ystate3d + 24,
    ystate3d + 32,  ystate3d + 40,  ystate3d + 48,   ystate3d + 56,
    ystate3d + 64,  ystate3d + 72,  ystate3d + 80,   ystate3d + 88,
    ystate3d + 96,  ystate3d + 104,  ystate3d + 112,  ystate3d + 120,
    ystate3d + 128,  ystate3d + 136,  ystate3d + 144,  ystate3d + 152,
    ystate3d + 160,  ystate3d + 168,  ystate3d + 176,  ystate3d + 184
};
```

Figure 14: State tables for Hilbert order

rules. Figure 15 is a C code for generating a Hilbert order in a three-dimensional space, which maps $[0, 1]^3$ to $[0, 1]$.

3.2.5 Space-filling Curve Partitioning Method

Space-filling curve methods are good alternatives for the graph methods, which can provide good partitioning. These methods are faster compared with the graph methods. Algorithm 2 shows the process of the space-filling curve methods, which has three steps. The first step is to map a computational domain Ω to a subset of $(0, 1)^3$, which can be obtained by a linear mapping. Then, for any cell, its new centroid coordinates belong to $(0, 1)^3$. A space-filling curve defines a map from $[0, 1]^3$ to $[0, 1]$. We can employ this map to calculate the values of the new centroid coordinates in $(0, 1)$. The third step is to partition $(0, 1)$ into N_p sub-intervals such that each sub-interval has the same numbers of cells (or workload). Each MPI can perform the first and the second steps independently. However, in our implementation, the third step is calculated by one MPI and it broadcasts results to other MPIs.

These methods assume that two cells (elements) that are close to each other have a higher

```

void HilbertInvOrder3d(HSFC_ENTRY *x)
{
    int level, EffLen;
    unsigned int key[3], c[3], temp, stat;
    INT i;

    static unsigned INTMX;
    static unsigned EfBit;
    static BOOLEAN initialized = FALSE;
    static int k0 = 0, k1 = 0, k2 = 0;

    if (!initialized) {
        initialized = TRUE;

        INTMX = 4294967295U;
        EfBit = INTMX >> 2;

        k0 = 60 - hsfc_maxlevel * 3;
        k1 = 30 - hsfc_maxlevel * 3;
        k2 = -hsfc_maxlevel * 3;
    }

    c[0] = (unsigned int)(x[i].coord[0] * (double)INTMX);
    c[1] = (unsigned int)(x[i].coord[1] * (double)INTMX);
    c[2] = (unsigned int)(x[i].coord[2] * (double)INTMX);
    c[1] >>= 1;
    c[2] >>= 2;

    key[0] = key[1] = key[2] = 0;
    stat = 0;
    EffLen = 30;
    for (level = 0; level < hsfc_maxlevel; level++) {
        EffLen--;
        temp = ((c[0] >> EffLen) & 4) | ((c[1] >> EffLen) & 2) | ((c[2] >> EffLen) & 1);

        key[0] = (key[0] << 3) | ((key[1] >> 27) & 7);
        key[1] = (key[1] << 3) | ((key[2] >> 27) & 7);
        key[2] = (key[2] << 3) | *(d[stat] + temp);

        stat = *(s[stat] + temp);
    }

    key[0] = key[0] & EfBit;
    key[1] = key[1] & EfBit;
    key[2] = key[2] & EfBit;

    x[i].hsfc = ldexp((double)key[2], k2);
    x[i].hsfc += ldexp((double)key[1], k1);
    x[i].hsfc += ldexp((double)key[0], k0);
}

```

Figure 15: Computation of Hilbert order

possibility to communicate with each other than two cells that are far from each other. This assumption is true for grid-based numerical methods, such as the finite element, finite volume and finite difference methods. The only difference for various space-filling curve methods is how to map

a cell to $(0, 1)$.

Algorithm 2 Space-filling curve method

- 1: Map a computational domain Ω to a subset of $(0, 1)^3$.
 - 2: For any cell, calculate its mapping that belongs to $(0, 1)$.
 - 3: Partition the interval $(0, 1)$ into N_p sub-intervals and each sub-interval has the same number of cells.
-

3.2.6 Partition Quality

This section studies partition quality on homogeneous architectures. The quality of a partition resulting from the graph methods and space-filling curve methods has several measurements, including a load imbalance factor, a local surface index and a global surface index. These concepts have been studied in [48].

The workload distribution is measured by the load imbalance factor. It is reasonable to assume that all cells have equal computations so the workload of each processor can be approximated by the number of cells it owns. The load imbalance factor f_p is defined as follows:

$$f_p = \frac{N_p \times \max_{1 \leq i \leq N_p} |\mathbb{G}_i|}{\sum_{i=1}^{N_p} |\mathbb{G}_i|}. \quad (9)$$

Let f_i be the number of faces in the sub-grid \mathbb{G}_i and b_i be the number of faces that are shared by a cell in another processor. We should mention that b_i is proportional to the communication volume of the i th processor involved.

The maximum local surface index is defined by

$$r_M = \max_{0 \leq i < n} \frac{b_i}{f_i}. \quad (10)$$

It is used to model the maximal communications that one processor involves. The global surface index is defined by

$$r_G = \frac{\sum_{i=1}^{N_p} b_i}{\sum_{i=1}^{N_p} f_i - \sum_{i=1}^{N_p} b_i}. \quad (11)$$

It models the overall communications of all processors. The average surface index is defined by

$$r_M = \sum_{i=1}^{N_p} \frac{b_i}{f_i}. \quad (12)$$

It is used to model the average communications that one processor involves.

If a cell in a sub-grid has a neighbour in another processor, then we can say that these two sub-grids are connected. Inter-processor connectivity is the number of connected sub-grids of a given sub-grid, denoted by c_i . The maximal inter-processor connectivity is defined by

$$c = \max_{1 \leq i \leq N_p} c_i. \quad (13)$$

```

typedef struct DOF_TYPE_
{
    INT          np_cell;      /* number of DOFs per cell */
    INT          np_well;      /* number of DOFs per well */
} DOF_TYPE;

typedef struct DOF_
{
    char         *name;        /* name of DOF */
    GRID         *g;           /* the grid */
    DOF_TYPE     *type;        /* type of DOF */
    FLOAT        *data;
    INT          *idata;
    INT          count_cell;   /* data count per cell */
    INT          count_well;   /* data count per well */
    INT          count_perf;   /* data count per perforation */
    INT          dim;
    BOOLEAN      assembled;
} DOF;

```

Figure 16: Data structure of DOF

4 DOF (Degrees of Freedom)

The cell-centered data is natural to reservoir simulation, since each cell can represent a block of a real oil and gas field and we can attach properties to the cell. Its data structure, DOF, is shown in Figure 16.

Each DOF has a name (**name**), and it associates with a grid (**g**). It also has a type (**type**). It can define floating-point number and integer data. A DOF also has a dimension (**dim**), through which scalar and vector can be defined. In reservoir simulations, cells have different properties, such as oil and water saturations, porosity and permeability.

```

DOF_TYPE DOF_CELL_      = {1, 0};
DOF_TYPE DOF_WELL_      = {0, 1};
DOF_TYPE DOF_PERF_      = {0, 1};
DOF_TYPE DOF_CONSTANT_  = {1, 0};

DOF_TYPE *DOF_CELL      = &DOF_CELL_;
DOF_TYPE *DOF_WELL      = &DOF_WELL_;
DOF_TYPE *DOF_PERF      = &DOF_PERF_;
DOF_TYPE *DOF_CONSTANT  = &DOF_CONSTANT_;

```

Figure 17: DOF types

DOF types are defined in Figure 17. DOF_TYPE_CELL defines data on each cell. DOF_TYPE_WELL defines data on each well. One well may perforate several cells and some properties only exist on these perforated cells, such as oil production rate and water injection rate. We have another type, DOF_TYPE_PERF, which has value on all perforated cells only. DOF_CONSTANT is for constant. The data types are floating-point number.

Cells may be in different MPIs, and sometimes, a cell requires data in its remote neighbouring

```
typedef struct DOF_NBR_
{
    DOF      *dof;
    FLOAT     *data;
    INT       *idata;
    INT       dim;
    BOOLEAN   assembled;
} DOF_NBR;
```

Figure 18: Remote DOF data

cell. `DOF_NBR` is defined to gather remote DOF data for cell's direct neighbours.

Input and output (reading data from a file and writing data to a file) for sequential applications are trivial, and simple writing and reading functions from operating systems or C language are enough. However, when we are working on parallel computing, each MPI task has portion of the whole grid and portion of data, and each task only reads and writes part of data. If each MPI task reads the whole input file and picks necessary data, then most read data will be dropped and they compete disk with each other. If all MPI tasks write data to the same file, conflicts may exist. The platform provides parallel input and output modules using MPI-IO, which supports integer and floating-point numbers.

The initialization of reservoir simulations requires read reservoir properties from files, such as porosity and permeability. The restart of simulations also needs to write intermediate data and to read data files. They are achieved by our parallel input and output modules.

5 MAP

```
typedef struct COMM_INFO_
{
    INT      *sidx;
    INT      *widx;
    int      *scnts, *sdsp;
    int      *rcnts, *rdsp;
    MPI_Comm comm;
    INT      ssize, rsize;
} COMM_INFO;
```

Figure 19: Data structure of communication info

Figure 19 shows communication info used by matrix and vector. This data structure can be used for collection communications, such as `MPI_Alltoallv`, and point-to-point communications. `scnts` means data amounts sent by current MPI task to other MPI tasks, and `sdsp` means the displacement relative to sending buffer. `rcnts` means data amounts received from other MPI tasks, and `rdsp` is the displacement relative to receiving buffer. `ssize` is total data sent by current MPI task and `rsize` is total received data. `sidx` is the index of data to be sent. `widx` is the local index of data received in current MPI task.

A map, `MAP`, is defined to store communication information among cell data, matrices and vectors. It includes DOF information, off-process entries, communication pattern, locations of data required by other MPI tasks and locations of data received from other MPI tasks.


```

typedef struct LG_MAP_
{
    INT *Gmap;
    INT *Lmap;
    INT size;
} LG_MAP;

typedef struct MAP_
{
    DOF          **_dofs;
    INT          ndofs;

    LG_MAP       *_RNGmap;
    INT          num_rnmaps;

    INT          *_L2Gmap;
    INT          *part;
    INT          *offset;

    INT          nlocal;
    INT          ntlocal;
    INT          nglobal;
    int          refcount;

    LG_MAP       *_W2V;
    INT          *_W2L;
    INT          num_w2vs;
    LG_MAP       **_P2V;
    INT          *num_perfs;
    INT          nwalls_global;

    MPI_Comm     comm;
    int          rank, nprocs;
    COMM_INFO    *cinfo;

    BOOLEAN      assembled;
} MAP;

```

Figure 20: Data structure for MAP

6 Distributed Matrices and Vectors

A linear system, $Ax = b$ ($A \in R^{N \times N}$), is assembled in each Newton iteration. Distributed matrix and vector are required to store the linear system. In the platform, each matrix and vector are distributed among all MPI tasks.

Each row of the distributed matrix has a unique global row index, which ranges from 0 to $N - 1$ consecutively and is numbered from the 1-th MPI task to the N_p -th MPI task. Each row also has a local index on each MPI task. The global indices of a vector is numbered the same way.

6.1 Vector

A distributed floating-point vector is defined in Figure 21, which has buffer that holds data entries (`data`), number of local entries belong to current MPI task (`nlocal`), number of total entries (including off-process entries, `ntlocal`) and reference to mapping information.

```
typedef struct VEC_  
{  
    MAP      *map;  
    FLOAT    *data;  
    INT      nlocal;    /* entries belong to current proc */  
    INT      ntlocal;   /* total entries in current MPI process */  
} VEC;
```

Figure 21: Data structure of VEC

6.2 Matrix

```
/* struct for a matrix row */  
typedef struct MAT_ROW_  
{  
    FLOAT    *data;     /* data */  
    INT      *cols;     /* local column indices, INT[ncols] */  
    INT      *gcols;    /* global column indices, INT[ncols] */  
    INT      ncols;     /* number of nonzero columns */  
} MAT_ROW;  
  
typedef struct MAT_  
{  
    MAT_ROW  *rows;  
  
    MAP      *map;  
    COMM_INFO *cinfo;  
    INT      *O2Gmap;  
  
    INT      nlocal;    /* local entries belong to current proc */  
    INT      ntlocal;   /* total local entries */  
    INT      nglobal;   /* global matrix size */  
    INT      *part;     /* distribution information */  
  
    int      rank;  
    int      nprocs;  
    MPI_Comm comm;  
    BOOLEAN  assembled;  
} MAT;
```

Figure 22: Data structure of MAT

The data structure of a distributed matrix is more complex than a vector, which requires entries for each row and some other additional information. It is represented in Figure 22. The `MAT_ROW`

stores non-zero entries of each row and its storage format is similar to a CSR matrix, which has a value of an entry (**data**), the global index (**gcol**) and local index of an entry (**col**). The **MAT** has communication information (**cinfo**), MPI information and additional information, such as mapping between local index of off-process entries and their global indices (**02Gmap**), and row distribution among all MPI tasks (**part**).

6.3 Algebraic Operations

With the help of above data structures, commonly used matrix-vector operations and vector operations can be implemented directly, which are listed as follows:

$$y = \alpha Ax + \beta y, \quad (14)$$

$$z = \alpha Ax + \beta y, \quad (15)$$

$$y = \alpha x + \beta y, \quad (16)$$

$$z = \alpha x + \beta y, \quad (17)$$

$$\alpha = \langle x, y \rangle, \quad (18)$$

$$\alpha = \langle x, x \rangle^{\frac{1}{2}}, \quad (19)$$

where A is a matrix, α and β are scalars, and x and y are vectors.

7 Linear System

For the linear system, $Ax = b$, derived from a nonlinear method, Krylov subspace solvers including the restarted GMRES(m) solver, the BiCGSTAB solver, and algebraic multi-grid (AMG) solvers are commonly used to find its solution. The Krylov subspace solvers mentioned here are suitable for arbitrary linear systems while the algebraic multi-grid solvers are efficient for positive definite linear systems.

7.1 Solvers

In-house parallel Krylov subspace linear solvers are implemented. The data structure of our solvers, **SOLVER**, is shown in Figure 23, which includes parameters (**rtol**, **atol**, **btol**, **maxit**, **restart**), matrices, right-hand sides, solutions, and preconditioner information.

7.2 Preconditioners

Several preconditioners are developed, including general purpose preconditioners and physics-based preconditioners for reservoir simulations only.

```

typedef struct SOLVER_
{
    FLOAT          residual;
    FLOAT          rtol;
    FLOAT          atol;
    FLOAT          btol;
    int            nits;
    INT            maxit;
    INT            restart;

    MAT            *A;
    VEC            *rhs;
    VEC            *x;
    MAP            *map;

    /* preconditioner */
    SOLVER_PC      pc;
    PC_TYPE        pc_type;

    int            rank;
    int            nprocs;
    MPI_Comm       comm;
    BOOLEAN        assembled;
    BOOLEAN        oem_created;
} SOLVER;

```

Figure 23: Data structure of SOLVER

7.2.1 Restricted Additive Schwarz Method

For the classical ILU methods, the given matrix A is factorized into a lower triangular matrix L and an upper triangular matrix U ; a lower triangular linear system and an upper triangular linear system are required to solve:

$$LUx = b \iff Ly = b, Ux = y. \quad (20)$$

The systems need to be solved row-by-row, which are serial. It is well-known that they have limited scalability. Another option for parallel computing is the restricted additive Schwarz (RAS) method [20], which was developed by Cai et al.

The data structure of the RAS preconditioner is shown in Figure 24. The `pars` stores parameters of the RAS preconditioner, such as overlap, local solver (ILUK, ILUT and ILTC), the level of ILUK, memory control and tolerance of ILUT, and filter tolerance. The `RAS_DATA` has a local problem stored by the lower triangular matrix L and the upper triangular matrix U , communication information of different sub-domains (`cinfo`), memory buffer and prolongation (restriction) operation information.

The default parameters of the RAS preconditioner is shown in Figure 25. Its default local solver is ILU(0). Default parameters for ILUT(p , tol) is -1 and 1e-3. If p is -1, the factorization subroutine will determine dynamically.

```

typedef struct RAS_PARS_
{
    INT      overlap;
    INT      iluk_level;
    INT      ilut_p;
    int      solver;
    FLOAT    ilut_tol;
    FLOAT    filter_tol;
    INT      ilutc_drop;
} RAS_PARS;

/* RAS */
typedef struct RAS_DATA_
{
    COMM_INFO      *cinfo;
    mat_csr_t      L;
    mat_csr_t      U;

    FLOAT          *frbuf;
    FLOAT          *fxbuf;
    FLOAT          *fbbuf;
    INT            *ras_pro;    /* ras prolongation */
    INT            num_ras_pro;

    RAS_PARS      pars;
} RAS_DATA;

```

Figure 24: Data structure of RAS preconditioner

```

static RAS_PARS ras_pars_default =
{
    /* overlap */      1,
    /* k */            0,
    /* ilut_p */       -1,
    /* solver */       ILUK,
    /* ilut_tol */     1e-3,
    /* filter tol */   1e-4,
    /* drop */         0,
};

```

Figure 25: Default parameters of RAS preconditioner

7.2.2 Algebraic Multigrid Methods

If A is a positive-definite square matrix, the AMG methods [31, 30, 27, 28, 29, 2] have proved to be efficient methods and they are also scalable [21]. AMG methods have hierarchical structures, and a coarse grid is chosen when entering a coarser level. Its structure of an algebraic multigrid solver is shown in Figure 26.

A restriction operator R_l and an interpolation (prolongation) operator P_l are determined. In general, the restriction operator R_l is the transpose of the interpolation (prolongation) operator P_l :

$$R_l = P_l^T.$$

The matrix on the coarser grid is calculated by

$$A_{l+1} = R_l A_l P_l. \quad (21)$$

We know that a high frequency error is easier to converge on a fine grid than a low frequency error, and for the AMG methods, the restriction operator, R_l , projects the error from a finer grid onto a coarser grid and converts a low frequency error to a high frequency error. The interpolation operator transfers a solution on a coarser grid to that on a finer grid. Its setup phase for the AMG methods on each level l ($0 \leq l < L$) is formulated in Algorithm 3, where a coarser grid, an interpolation operator, a restriction operator, a coarser matrix and post- and pre-smoothers are constructed. By repeating the algorithm, an L -level system can be built. The solution of the AMG methods is recursive and is formulated in Algorithm 4, which shows one iteration of AMG. The AMG package we use is the BoomerAMG from HYPRE [33].

The Cleary-Luby-Jones-Plassman (CLJP) parallel coarsening algorithm was proposed by Cleary [23] based on the algorithms developed by Luby [26] and Jones and Plassman [25]. The standard RS coarsening algorithm has also been parallelized [33]. Falgout et al. developed a parallel coarsening algorithm, the Falgout coarsening algorithm, which has been implemented in HYPRE [33]. Yang et al. proposed HMIS and PMIS coarsening algorithms for a coarse grid selection [1]. Various parallel smoothers and interpolation operators have also been studied by Yang et al [33, 1].




Figure 26: Structure of AMG solver.

Algorithm 3 AMG setup Algorithm

- 1: Calculate strength matrix S .
 - 2: Choose coarsening nodes set ω_{l+1} according to strength matrix S , such that $\omega_{l+1} \subset \omega_l$.
 - 3: Calculate prolongation operator P_l .
 - 4: Derive restriction operator $R_l = P_l^T$.
 - 5: Calculate coarse matrix A_{l+1} : $A_{l+1} = R_l \times A_l \times P_l$.
 - 6: Setup pre-smoother S_l and post-smoother T_l .
-

In our platform, the AMG solvers are from HYPRE. They can be used as solvers and preconditioners. The data structure of the AMG solvers is shown in Figure 27. The `BMAMG_PARS` stores

Algorithm 4 AMG V-cycle Solution Algorithm: $\text{amg_solve}(l)$

Require: $b_l, x_l, A_l, R_l, P_l, S_l, T_l, 0 \leq l < L$

```
 $b_0 = b$ 
if ( $l < L$ ) then
     $x_l = S_l(x_l, A_l, b_l)$ 
     $r = b_l - A_l x_l$ 
     $b_{l+1} = R_l r$ 
     $\text{amg\_solve}(l + 1)$ 
     $x_l = x_l + P_l x_{l+1}$ 
     $x_l = T_l(x_l, A_l, b_l)$ 
else
     $x_l = A_l^{-1} b_l$ 
end if
 $x = x_0$ 
```

```
typedef struct BMAMG_PARS_
{
    INT          maxit;
    INT          num_funcs;           /* size of the system of PDEs */
    INT          max_levels;         /* max MG levels */
    FLOAT        strength;           /* strength threshold */
    FLOAT        max_row_sum;        /* max row sum */
    FLOAT        trunc_tol;          /* trunc tol */

    int          coarsen_type;        /* default coarsening type = Falgout */
    int          cycle_type;          /* MG cycle type */
    int          relax_type;          /* relaxation type */
    int          coarsest_relax_type; /* relax type on the coarsest grid */
    int          interp_type;         /* interpolation */
    INT          num_relax;           /* number of sweep */

} BMAMG_PARS;

typedef struct BMAMG_DATA_
{
    MAP          *map;
    INT          ilower, iupper;
    HYPRE_IJMatrix A;
    HYPRE_IJVector b, x;
    BOOLEAN      assembled;
    HYPRE_Solver  hsolver;

    SolveFcn     setup;
    SolveFcn     solve;
    DestroyFcn    destroy;

    BMAMG_PARS   pars;

} BMAMG_DATA;
```

Figure 27: Data structure of AMG method

```

static BMAMG_PARS amg_pars_default =
{
    /* maxit */           1,
    /* num_funcs */       -1,
    /* max_levels */      6,
    /* strength */        0.5,
    /* max_row_sum */     0.9,
    /* trunc error */     1e-2,

    /* coarsen_type */    Falgout,
    /* cycle_type */      v-cycle,
    /* relax_type */      gs-h-forward,
    /* coarsest_relax_type */ gs-h-symmetric,
    /* interp type */     cmi,
    /* itr relax */       2,
};

```

Figure 28: Default parameters of AMG method

parameters of the AMG method, including the coarsening type, interpolation type, maximal levels, smoother type, and cycle type. The `BMAMG_DATA` stores linear system information, such as a distribution pattern of matrices and vectors, mapping information, and related interfaces.

Default parameters for the AMG method is shown in Figure 28, where a default six-level AMG method is applied. The detailed explanation of each parameter can be read from the HYPRE user manual.

7.2.3 CPR-like Preconditioners

Linear systems from black oil and composition models are hard to solve, especially when the reservoirs are heterogeneous. However, the matrices from the pressure unknowns are positive definite, which can be solved by AMG methods. Many preconditioners have been developed to speed the solution of linear systems, such as the constrained pressure residual (CPR) method and FASP method [18]. Here we introduce our multi-stage preconditioners for the black oil and compositional models, which are based on the classical CPR method.

Numerical methods for black oil and compositional models may choose different unknowns [32]. Here we assume that the oil phase pressure p_o is always one of the unknowns. The other variables are denoted as \vec{s}_i . The well unknowns are denoted by \vec{w} , whose dimension equals the number of wells in the reservoir, N_w . Let us define the pressure vector p as:

$$p = \begin{pmatrix} p_{o,1} \\ p_{o,2} \\ \dots \\ p_{o,N_g} \end{pmatrix}, \quad (22)$$

and the global unknown vector x as:

$$x = \begin{pmatrix} p_{o,1} \\ p_{o,2} \\ \dots \\ p_{o,N_g} \\ \overrightarrow{s_1} \\ \dots \\ \overrightarrow{s_{N_g}} \\ \overrightarrow{w} \end{pmatrix}. \quad (23)$$

A restriction operator from x to p is defined as

$$\Pi_r x = p. \quad (24)$$

A prolongation operator Π_p is defined as

$$\Pi_p p = \begin{pmatrix} p \\ 0 \end{pmatrix}, \quad (25)$$

where $\Pi_p p$ has the same dimension as x .

If a proper ordering technique is applied, the matrix A from black oil and compositional models can be written as equation (26),

$$A = \begin{pmatrix} A_{pp} & A_{ps} & A_{pw} \\ A_{sp} & A_{ss} & A_{sw} \\ A_{wp} & A_{ws} & A_{ww} \end{pmatrix}, \quad (26)$$

where the sub-matrix A_{pp} is the matrix corresponding to the pressure unknowns, the sub-matrix A_{ss} is the matrix corresponding to the other unknowns, the sub-matrix A_{ww} is the matrix corresponding to the well bottom hole pressure unknowns, and other matrices are coupled items.

Let us introduce some notations for the preconditioning system $Ax = f$. If A is a positive definite matrix, then we define the notation $\mathcal{M}_g(A)^{-1}b$ to represent the solution x from AMG methods. If it is solved by the RAS method, then we use the notation $\mathcal{R}(A)^{-1}b$ to represent solution x . The CPR-like preconditioners we develop are shown by Algorithm 5 to Algorithm 8, which are noted as CPR-FP, CPR-PF, CPR-FPF and CPR-FFPF methods [40], respectively.

Algorithm 5 The CPR-FP preconditioner

- 1: $x = \mathcal{R}(A)^{-1}f$.
 - 2: $r = f - Ax$
 - 3: $x = x + \Pi_p(\mathcal{M}_g(A_{pp})^{-1}(\Pi_r r))$.
-

The data structure of the CPR preconditioners is shown in Figure 29. It has a RAS solver (**ras**) an AMG solver (**amg**), prolongation information (**pro_pres** and **num_pro_pres**), buffers and settings for RAS solver and AMG solver. The term **CPR_PARS** stores parameters of the CPR methods.

Algorithm 6 The CPR-PF preconditioner

- 1: $x = \Pi_p(\mathcal{M}_g(A_{pp})^{-1}(\Pi_r f))$
 - 2: $x = \mathcal{R}(A)^{-1}f$.
 - 3: $r = f - Ax$
 - 4: $x = x + \mathcal{R}(A)^{-1}f$.
-

Algorithm 7 The CPR-FPF preconditioner

- 1: $x = \mathcal{R}(A)^{-1}f$.
 - 2: $r = f - Ax$
 - 3: $x = x + \Pi_p(\mathcal{M}_g(A_{pp})^{-1}(\Pi_r r))$.
 - 4: $r = f - Ax$
 - 5: $x = x + \mathcal{R}(A)^{-1}r$.
-

Algorithm 8 The CPR-FFPF preconditioner

- 1: $x = \mathcal{R}(A)^{-1}f$.
 - 2: $r = f - Ax$
 - 3: $x = x + \mathcal{R}(A)^{-1}r$.
 - 4: $r = f - Ax$
 - 5: $x = x + \Pi_p(\mathcal{M}_g(A_{pp})^{-1}(\Pi_r r))$.
 - 6: $r = f - Ax$
 - 7: $x = x + \mathcal{R}(A)^{-1}r$.
-

```
typedef struct CPR_PARS_  
{  
    RAS_PARS    ras;  
    BMAMG_PARS  amg;  
    INT         pres_which;  
    INT         pres_loc;  
  
    INT         itr_ras_pre;  
    INT         itr_ras_post;  
  
} CPR_PARS;  
  
typedef struct CPR_DATA_  
{  
    RAS_DATA     ras;  
    BMAMG_DATA   amg;  
  
    INT          *pro_pres;    /* prolongation from pressure */  
    INT          num_pro_pres;  
    VEC          *varbuf;      /* vector r (A), buffer */  
    VEC          *vpbbuf;      /* buffer for AMG */  
    VEC          *vpxbuf;      /* buffer for AMG */  
  
    CPR_PARS     pars;  
  
} CPR_DATA;
```

Figure 29: Data structure of CPR preconditioners

```

/* pc type */
typedef enum PC_TYPE_
{
    PC_RAS,           /* Restricted Additive Schwarz */
    PC_AMG,           /* BommerAMG */
    PC_CPR_FP,        /* cpr */
    PC_CPR_PF,        /* cpr */
    PC_CPR_FPF,       /* cpr */
    PC_CPR_FFPF,      /* cpr */
    PC_USER,          /* user define */
    PC_NON,           /* no preconditioner */
} PC_TYPE;

/* preconditioner interface */
typedef void (*PC_USER_ASSEMBLE)(struct SOLVER_PC_ *pc, MAT *mat);
typedef void (*PC_SOLVE)(struct SOLVER_PC_ *pc, VEC *x, VEC *b);
typedef void (*PC_DESTROY)(struct SOLVER_PC_ *pc);

/* SOLVER_PC struct */
typedef struct SOLVER_PC_
{
    struct SOLVER_      *solver;      /* pointer to solver */

    void                *data;
    PC_USER_ASSEMBLE    user_assemble;

    /* MPI */
    int                 rank;
    int                 nprocs;
    MPI_Comm            comm;

    PC_SOLVE            solve;
    PC_DESTROY          destroy;
    BOOLEAN             assembled;
} SOLVER_PC;

```

Figure 30: Data structure of preconditioners

7.3 Data Structure for Preconditioners

The data structure for preconditioners is defined by Figure 30. It provides three function pointers that can complete assembling (PC_USER_ASSEMBLE), solving (PC_SOLVE) and destroying (PC_DESTROY) a preconditioning system. With these function pointers, this data structure is general purpose, and if a new set of implementations are provided, a new preconditioner can be constructed. From the data structure, we can see built-in preconditioners, including RAS method, AMG methods and CPR methods, are provided, and users can implement their own preconditioners by providing proper assembling, solving and destroying functions. The data structure also contains a pointer to solver, and MPI info.

8 Numerical Experiments

A Blue Gene/Q from IBM that located in the IBM Thomas J. Watson Research Center is employed. The system uses PowerPC A2 processor. Each processor has 18 cores and 16 cores are used for computation. Performance of each core is really low compared with processors from Intel. However, it has a strong network relative to compute performance and the system is scalable. Since the platform is designed for parallel applications, scalability is the most important objective. In the following section, we will focus on scalability.

8.1 Grid Partitioning

Example 1 *Two grids are applied to test the partitioning quality of the Hilbert space-filling curve method (HSFC), compared with other methods, such as the HSFC method from Zoltan, RCB (recursive coordinate bisection) from Zoltan, the Morton space-filling curve method, and ParMETIS. The first grid is Ω_1 , which is a long cylinder and contains 2,522,624 cells. The second one is a thin plate with many holes, which has 3,713,792 cells. They are shown in Figures 31 and 32. Numerical results are presented in Tables 1 and 2 [54]. These grids are used for the finite element method and numerical experiments are from PHG (Parallel Hierarchical Grid) [11].*



Figure 31: Grid Ω_1

Table 1: Maximal inter-processor connectivities: Ω_1 and Ω_2

Ω_1						
# submeshes	16	32	64	128	160	192
ParMETIS	2	2	5	8	8	9
RCB	2	2	5	11	12	13
MSFC	4	5	10	16	19	21
HSFC	3	6	13	23	24	23
Zoltan/HSFC	12	18	21	23	24	24
Ω_2						
# submeshes	16	32	64	128	160	192
ParMETIS	7	7	7	10	11	12
RCB	7	7	8	11	12	12
MSFC	9	13	18	21	22	25
HSFC	8	10	13	18	20	21
Zoltan/HSFC	12	19	23	25	27	31

ParMETIS is a package for a graph method and other methods are geometric methods. In Table 1, the inter-processor connectivity models the number of startups in group communication, which measures latency. A smaller index means less latency, which is good for parallel computing. We




Figure 32: Grid Ω_2

Table 2: Surface indices for Ω_1

maximum surface index ($\%$, Ω_1)						
# submeshes	16	32	64	128	160	192
ParMETIS	0.74	1.48	2.88	5.00	5.78	5.71
RCB	1.0	2.18	4.07	6.06	7.06	7.60
MSFC	4.2	6.51	10.6	16.2	18.9	20.7
HSFC	3.80	8.31	15.5	19.5	20.4	21.5
Zoltan/HSFC	13.9	19.1	25.4	28.1	30.9	35.8
average surface index ($\%$, Ω_1)						
# submeshes	16	32	64	128	160	192
ParMETIS	0.64	1.28	2.39	3.62	4.05	4.42
RCB	0.85	1.68	3.43	5.17	6.03	6.58
MSFC	2.93	5.14	7.44	9.8	10.7	11.5
HSFC	2.78	5.03	7.18	9.34	10.2	10.8
Zoltan/HSFC	8.72	11.8	16.2	20.6	22.0	23.6

can see that ParMETIS is the best and the RCB method has similar quality as ParMETIS. Other methods are worse than these two methods.

Tables 2 and 3 present surface indices, which model a communication volume. A smaller value means less communication. These two tables demonstrate again that the graph methods can minimize communication. The results show that the RCB method has better partitioning quality than other space-filling curve methods. However, our HSFC method has better quality than the HSFC

Table 3: Surface indices for Ω_2						
maximum surface index ($\%$, Ω_2)						
# submeshes	16	32	64	128	160	192
ParMETIS	2.45	2.56	3.30	5.32	6.13	6.10
RCB	2.57	4.74	6.29	9.16	9.66	11.6
MSFC	3.0	5.11	7.06	9.74	11.1	11.8
HSFC	2.86	4.74	7.03	9.78	10.6	11.3
Zoltan/HSFC	8.89	12.3	16.5	19.5	21.2	22.8
average surface index ($\%$, Ω_2)						
# submeshes	16	32	64	128	160	192
ParMETIS	1.12	1.45	2.18	3.43	3.87	4.21
RCB	1.89	3.04	3.93	5.90	6.58	7.2
MSFC	2.24	3.61	5.21	7.34	8.22	8.92
HSFC	2.19	3.29	4.88	6.92	7.81	8.27
Zoltan/HSFC	5.10	6.86	9.47	12.6	14.1	15.3

method from the Zoltan package.

Example 2 *The computational domain is $1200\text{ft} \times 2200\text{ft} \times 170\text{ft}$ and the grid size is $180 \times 660 \times 255$, which has 30,294,000 cells. The HSFC method and ParMETIS are applied to test partitioning quality. Numerical results are shown in Tables 4 and 5.*

Table 4: Maximal inter-processor connectivities, Example 2

Example 2				
# submeshes	256	512	1024	2048
ParMETIS	19	20	21	26
HSFC	16	15	15	16

Table 5: Surface indices for Eaxmple 2

maximum surface index ($\%$)				
# submeshes	256	512	1024	2048
ParMETIS	6.86	9.23	11.0	13.6
HSFC	8.34	10.7	13.5	18.2
average surface index ($\%$)				
# submeshes	256	512	1024	2048
ParMETIS	5.15	6.55	8.37	10.5
HSFC	6.39	8.16	10.4	13.0

Table 4 shows inter-processor connectivity. We can see that in this case the HSFC method has better inter-processor connectivity, and the resulting partitioning has less latency. Table 5 shows that partitioning from the HSFC method has more communication than ParMETIS.

Example 3 This example solves the *Helmholtz* problem using an adaptive finite element method, which has the *Dirichlet* boundary:

$$\begin{cases} -\Delta u + u = f, & (x, y, z) \in \Omega, \\ u(x, y, z) = g & \text{on } \partial\Omega. \end{cases}$$

The computation domain is Ω_1 , and the equation has the analytical solution:

$$u = \cos(2\pi x) \cos(2\pi y) \cos(2\pi z).$$

The initial grid has 4,927 cells (elements). The error indicator is defined by

$$\eta_K^2 = h_K^2 \|f_h + \Delta u_h + u_h\|_{L^2(K)}^2 + \sum_{f \subset \partial K \cap \Omega} \frac{h_f}{2} \left\| \left[\frac{\partial u_h}{\partial n_f} \right] \right\|_{L^2(f)}^2.$$

Several partitioning methods are compared, including the HSFC (PHG/HSFC) method, the HSFC method from Zoltan (Zoltan/HSFC), RCB, ParMETIS, and the refinement-tree method (RTK).




Figure 33: Partitioning time, Example 3

This example studies partitioning time and dynamic load balancing time. The dynamic load balancing includes several stages, such as grid partitioning, data (DOF, degrees of freedom) exchange, and sub-grid assembly. An adaptive finite element method is applied and the final grid size is around 12 million cells (elements). Figure 33 shows that the refinement tree method is the fastest method. The space-filling curve methods have a similar speed. The ParMETIS is slower than these space-filling curve methods, but its speed is similar to the RCB method. Figure 34 shows that the




Figure 34: Dynamic load balancing time, Example 3

RTK method uses the least dynamic load balancing time and is the fastest method, followed by ParMETIS. The Morton space-filling curve method has a similar speed to our HSFC method. The RCB method is slightly slower but it is faster than the HSFC method from the Zoltan package.

Example 4 *This case tests a two-phase oil-water model with a refined SPE10 project, where each original cell is refined into 27 smaller cells. The model has around 30 millions of cells. The nonlinear method is an inexact Newton method. The solver is GMRES(50) and the preconditioner is the CPR-PF method [40]. Only five time steps are applied. The case is run on GPC. Numerical summaries are shown in Table 6.*

The CPR-PF method is a combination of the algebraic multi-grid method (AMG) and the restricted additive Schwarz method (RAS) [22]. The grid partitioning affects the distribution of matrices and these two numerical methods. Table 6 shows that the number of Newton iterations varies from 79 to 106 under different settings. The linear iterations are from 231 to 379. The overall running time shows that when using 256 and 512 MPIs, the simulations with the HSFC method are faster, and when using 1024 and 2048 MPIs, the simulations with ParMETIS are faster. Considering the average running time for each Newton iteration, the simulations with ParMETIS use less time than simulations with the HSFC method.

Example 5 *This case tests a three-phase black oil problem with a refined SPE10 geological model, where each cell is refined into 27 smaller cells. The model has around 30 millions of cells. The nonlinear method is an inexact Newton method. The solver is GMRES(30) and the preconditioner is the CPR-PF method [40]. Only five time steps are applied. The case is run on GPC again. Numerical summaries are shown in Table 7.*

Table 6: Numerical summaries for Eaxmple 4

# Newton iterations				
# MPIs	256	512	1024	2048
ParMETIS	106	105	105	79
HSFC	101	80	105	101
# Linear iterations				
# MPIs	256	512	1024	2048
ParMETIS	364	231	332	375
HSFC	251	289	379	296
Overall running time (s)				
# MPIs	256	512	1024	2048
ParMETIS	668.5	346.6	236.1	116.4
HSFC	621.0	288.4	295.2	187.2
Average time for each Newton iteration (s)				
# MPIs	256	512	1024	2048
ParMETIS	6.30	3.30	2.25	1.47
HSFC	6.15	3.61	2.81	1.85

Table 7: Numerical summaries for Eaxmple 5

# Newton iterations				
# MPIs	256	512	1024	2048
ParMETIS	22	23	19	23
HSFC	22	22	22	22
# Linear iterations				
# MPIs	256	512	1024	2048
ParMETIS	141	145	110	149
HSFC	146	132	142	142
Overall running time (s)				
# MPIs	256	512	1024	2048
ParMETIS	456.8	259.2	138.3	249.6
HSFC	441.4	234.3	143.4	118.8
Average time for each Newton iteration (s)				
# MPIs	256	512	1024	2048
ParMETIS	20.7	11.3	7.28	10.85
HSFC	20.1	10.7	6.52	5.4

This example shows different behavior from Example 4. Table 7 shows that the simulations with the HSFC partitioning method are faster than the simulations with ParMETIS.

8.2 SpMV

Example 6 *This example tests the performance of sparse matrix-vector multiplication (SpMV) on IBM Blue Gene/Q. The matrix is a square matrix from a Poisson equation and it has 200 millions rows. Its performance is shown in Table 8 and its scalability is presented in Figure 35.*

# procs	32	64	128	256	512	1024
Time (s)	2.211	1.078	0.556	0.269	0.134	0.067

Table 8: Performance of SpMV, Example 6




Figure 35: Scalability of SpMV, Example 6

This example uses up to 128 compute cards, when more than 128 MPI tasks are used, multiple MPI tasks run on one card. From Table 8, we can see that when MPI tasks are doubled, the running time of SpMV is reduced by half. This example show our SpMV kernel has excellent scalability. Speedup is compared with ideal condition in Figure 35, which shows that our SpMV kernel has good scalability.

8.3 Poisson Equation

Example 7 *This example tests a Poisson equation with 3 billions of grid cells. The example uses up to 4,096 CPU cores (MPIs). The linear solver is GMRES(30) method with RAS preconditioner. The solver runs 90 iterations. The overlap of RAS method is one, and sub-domain problem on each core is solved by ILU(0). The numerical summaries are reported in Table 9 and scalability results are shown in Figure 36.*

# procs	Gridding	Building	Assemble	Overall (s)	Speedup
512	217.0	29.16	66.42	918.91	1.0
1024	98.83	14.79	33.71	454.04	2.02
2048	47.53	7.49	17.47	227.05	4.05
4096	23.31	3.86	9.17	116.64	7.88

Table 9: Numerical summaries of Example 7




Figure 36: Scalability of Example 7

Table 9 shows numerical summaries of Poisson equation. This example tests strong scalability of our platform. Here the gridding time includes grid generation, grid partitioning and grid redistribution. From Table 9, we can see that the gridding has excellent scalability and when MPIs are doubled, running time for gridding reduces by half. Building time is the time spent on generation of a linear system $Ax = b$. Since there is no communication involved, the scalability is ideal. Assemble time includes time for generating linear solver, and time for generating preconditioner (RAS method). The overall time includes gridding time, building time, assemble time and solution time. Again, from Table 9 and Figure 36, we can see our platform has excellent scalability.

Example 8 *The case goes with the size of the matrix at 6 billion variables and is constructed from pressure equations. GMRES(30) was applied to solve the system with RAS (Restricted Additive Schwarz) as the preconditioner and fixed iterations at 90. IBM Blue Gene/Q was used to carry the simulation. Numerical summaries are shown in Table 10 and scalability is presented in Fig 37.*

Table 10: Numerical summaries of the large-scale model

# procs	Gridding (s)	Build (s)	Assemble (s)	Overall (s)
512	429.78	57.83	130.54	1829.18
1024	200.24	29.28	66.12	906.73
2048	106.83	14.83	34.03	463.59
4096	47.82	7.59	17.7	232.77

This example tests the scalability of grid generation, building of linear system, and solution of linear system (including solver, preconditioner and SpMV). Table 10 shows that when MPI tasks are doubled, running time of grid generation, building of linear system, and solution of linear system is cut by half, which means our platform and linear solvers have excellent scalability. The scalability




Figure 37: Scalability of the large-scale model

is demonstrated by Figure 37 demonstrate. This example also show that the solver can calculate extremely large-scale linear systems.

8.4 Oil-water Model

Example 9 *This example tests a refined SPE10 case for the two-phase oil-water model, where each grid cell is refined into 27 grid cells. This case has around 30 millions of grid cells and around 60 millions of unknowns. The stopping criterion for the inexact Newton method is $1e-3$ and the maximal Newton iterations are 20. The BiCGSTAB solver is applied and its maximal iterations are 100. The preconditioner is the CPR-FPF preconditioner. The potential reordering and the Quasi-IMPES decoupling strategy are applied. The simulation period is 10 days. Up to 128 compute cards are used. The numerical summaries are shown in Table 11, and the speedup (scalability) is shown in Figure 38.*

# procs	# Steps	# Ntn	# Slv	# Avg-S	Time (s)	Avg-T (s)
64	50	315	3451	10.9	119167.4	378.3
128	48	286	3296	11.5	49488.7	173.0
256	54	323	4190	12.9	30423.2	94.1
512	52	329	3635	11.0	14276.5	43.3
1024	54	316	3969	12.5	7643.9	24.1

Table 11: Numerical summaries of Example 9

In this example, up to 1,024 MPI tasks are employed and the simulation with 64 MPI tasks is used as the base case to calculate speedup and scalability. The numerical summaries in Table 11 show the inexact Newton method is robust, where around 50 time steps and around 300 Newton




Figure 38: Scalability of Example 9

iterations are used for each simulation with different MPI tasks. The linear solver BiCGSTAB and the preconditioner CPR-FPF show good convergence, where the average number of linear iterations for each nonlinear iteration is between 10 and 13. The results mean our linear solver and preconditioner are effective and efficient. The overall running time and average time for each Newton iteration show our simulator has excellent scalability on IBM Blue Gene/Q, which is almost ideal for parallel computing. The scalability is also demonstrated in Figure 38. The running time and scalability curve also demonstrate our linear solver and preconditioner are scalable for large-scale simulation.

Example 10 *This example tests a refined SPE10 case for the two-phase oil-water model, where each grid cell is refined into 125 grid cells. This case has around 140 millions of grid cells and around 280 millions of unknowns. The stopping criterion for the inexact Newton method is $1e-2$ and the maximal Newton iterations are 20. The GMRES(30) solver is applied and its maximal iterations are 100. The preconditioner is the CPR-FPF preconditioner. The potential reordering and the Quasi-IMPES decoupling strategy are applied. The simulation period is 2 days. Up to 128 compute cards are used. The numerical summaries are shown in Table 12, and the speedup (scalability) is shown in Figure 39.*

# procs	# Steps	# Ntn	# Slv	# Avg-S	Time (s)	Avg-T (s)
256	36	225	5544	24.6	117288.1	521.2
512	36	226	5724	25.3	57643.1	255.0
1024	35	207	5446	26.3	27370.3	132.2
2048	36	209	5530	26.4	14274.9	68.3

Table 12: Numerical summaries of Example 10




Figure 39: Scalability of Example 10

In this example, up to 2048 MPI tasks are employed and the simulation with 256 MPI tasks is used as the base case to calculate speedup and scalability. The numerical summaries in Table 12 show the inexact Newton method is robust, where around 36 time steps and around 220 Newton iterations are used for each simulation with different MPI tasks. The linear solver GMRES(30) and the preconditioner CPR-FPF show good convergence, where the average number of linear iterations for each nonlinear iteration is around 26. The overall running time and average time for each Newton iteration show our simulator has excellent scalability on IBM Blue Gene/Q, which is almost ideal for parallel computing and shows slight super-linear scalability. The scalability is also demonstrated in Figure 39. The results show our linear solver and preconditioner are scalable for large-scale simulation.

8.5 Black Oil Model

Example 11 *The example tests the scalability of the black oil simulator using a refined SPE10 geological model, where each grid cell is refined to 27 grid cells. The model has 30.3 millions of grid cells. The inexact Newton method is applied and the termination tolerance is 10^{-2} . The linear solver is BiCGSTAB, whose maximal inner iterations are 100. The preconditioner is the CPR-FPF method and the overlap for the RAS method is one. The potential reordering and the ABF methods are enabled. The simulation period is 10 days. The maximal change allowed in one time step of pressure is 1,000 psi and the maximal change of saturation is 0.2. Up to 128 compute cards are used. Summaries of numerical results are shown in Table 13.*

Table 13 includes information for the nonlinear method, linear solver and running time. For all simulations, 33 time steps are used and the total Newton iterations are around 300. The results show the inexact Newton method is robust. For the linear solver and preconditioner, their convergence is good, which terminate in around 4 iterations. The results mean the linear solver and preconditioner are robust and effective for this highly heterogeneous model. The running time, average time per

# procs	# Steps	# Ntn	# Slv	# Avg-S	Time (s)
64	33	292	1185	4.0	106265.9
128	33	296	1150	3.8	50148.3
256	33	299	1267	4.2	25395.8
512	33	301	1149	3.8	12720.5
1024	33	301	1145	3.8	6814.2

Table 13: Numerical summaries of Example 11




Figure 40: Scalability (speedup) of Example 11

Newton iteration and scalability curve in Figure 40 show the scalability of our simulator, linear solver and preconditioner is good. When we use up to 1,024 MPI tasks and each compute card runs up to 8 MPI tasks, the scalability is excellent.

Example 12 *The case is a refined SPE1 project with 100 millions of grid cells. Linear solver is BiCGSTAB. Potential reordering and ABF decoupling are applied. Numerical summaries are in Table 14.*

MPIs	# Steps	# Newton	# Solver	# Avg. Itr	Time (s)
512	27	140	586	4.1	11827.9
1024	27	129	377	2.9	5328.4
2048	26	122	362	2.9	2708.5
4096	27	129	394	3.0	1474.2

Table 14: Numerical summaries, Example 12




Figure 41: Scalability of Example 12

9 Conclusion

Our work on developing an in-house parallel platform is presented in this paper, which provides grids, data, linear solvers and preconditioners for reservoir simulators. Various techniques and methods have been introduced, including the Hilbert space-filling curve method, topological partitioning method, structured grid, distributed matrices and vectors, and multi-state preconditioners for reservoir simulations. Examples, including grid management, linear solvers, pressure equations, and the black oil simulator, are presented to benchmark our platform. Numerical results show that our platform and simulators have excellent scalability and applications based on the platform can be sped up thousands of times faster. This paper also shows parallel computing is a powerful tool for large-scale scientific computing.

Acknowledgement

The support of Department of Chemical and Petroleum Engineering and Reservoir Simulation Group, University of Calgary is gratefully acknowledged. The research is partly supported by NSERC/AIEES/Foundation CMG and AITF Chairs.

References

- [1] Hans DS, Yang UM, Heys J, Reducing Complexity in Parallel Algebraic Multigrid Preconditioners, SIAM Journal on Matrix Analysis and Applications 27, (2006), 1019-1039.
- [2] UM Yang, On the Use of Relaxation Parameters in Hybrid Smoothers, Numerical Linear Algebra With Applications, 11, (2004), 155-172.

- [3] J. E. Killough and R. Bhogeswara. Simulation of compositional reservoir phenomena on a distributed-memory parallel computer. *Journal of Petroleum Technology* 43.11 (1991): 1368-1374.
- [4] J. M. Rutledge, D. R. Jones, W. H. Chen, and E. Y. Chung, The Use of Massively Parallel SIMD Computer for Reservoir Simulation, SPE-21213, eleventh SPE Symposium on Reservoir Simulation, Anaheim, 1991.
- [5] G. Shiralkar, R.E. Stephenson, W. Joubert, O. Lubeck, and B. van Bloemen Waanders, A production quality distributed memory reservoir simulator, SPE Reservoir Simulation Symposium. 1997.
- [6] T. Kaarstad, J. Froyen, P. Bjorstad, M. Espedal, Massively Parallel Reservoir Simulator, SPE-29139, presented at the 1995 Symposium on Reservoir Simulation, San Antonio, Texas, 1995.
- [7] J. E. Killough, D. Camilleri, B.L. Darlow, J. A. Foster, Parallel Reservoir Simulator Based on Local Grid Refinement, SPE-37978, SPE Reservoir Simulation Symposium, Dallas, 1997.
- [8] A.H. Dogru, H.A. Sunaidi, L.S. Fung, W.A. Habiballah, N. Al-Zamel, K.G. Li, A parallel reservoir simulator for large-scale reservoir simulation, SPE Reservoir Evaluation & Engineering 5.1 (2002): 11-23.
- [9] A.H. Dogru, L. S. Fung, U. Middy, T. Al-Shaalan, J.A. Pita, A next-generation parallel reservoir simulator for giant reservoirs, SPE/EAGE Reservoir Characterization & Simulation Conference. 2009.
- [10] L. Zhang, A Parallel Algorithm for Adaptive Local Refinement of Tetrahedral Meshes Using Bisection, *Numer. Math.: Theory, Methods and Applications*, 2009, 2, 65–89.
- [11] L. Zhang, T. Cui, and H. Liu, A set of symmetric quadrature rules on triangles and tetrahedra, *J. Comput. Math*, 2009, 27(1), 89–96.
- [12] Y. Saad, *Iterative methods for sparse linear systems*, SIAM, 2003.
- [13] J.R. Wallis, Incomplete Gaussian elimination as a preconditioning for generalized conjugate gradient acceleration, SPE Reservoir Simulation Symposium, 1983.
- [14] J.R. Wallis, R. P. Kendall, and T. E. Little, Constrained residual acceleration of conjugate residual methods, SPE Reservoir Simulation Symposium, 1985.
- [15] H. Cao, T. Schlumberger, A. Hamdi, J.R. Wallis, H.E. Yardumian, Parallel scalable unstructured CPR-type linear solver for reservoir simulation. SPE Annual Technical Conference and Exhibition. 2005.
- [16] R. E. Bank, T. F. Chan, W. M. Coughran Jr., R. K. Smith, The Alternate-Block-Factorization procedure for systems of partial differential equations, *BIT Numerical Mathematics* 29.4 (1989): 938-954.
- [17] T. M. Al-Shaalan, H. M. Klie, A. H. Dogru, M. F. Wheeler, Studies of Robust Two Stage Preconditioners for the Solution of Fully Implicit Multiphase Flow Problems. SPE Reservoir Simulation Symposium. 2009.

- [18] X. Hu, W. Liu, G. Qin, J. Xu, Z. Zhang, Development of a fast auxiliary subspace preconditioner for numerical reservoir simulators, SPE Reservoir Characterisation and Simulation Conference and Exhibition. 2011.
- [19] T. Chen, N. Gewecke, Z. Li, A. Rubiano, R. Shuttleworth, B. Yang and X. Zhong, Fast Computational Methods for Reservoir Flow Models, Technical report, University of Minnesota, 2009.
- [20] A. Elli, and O. B. Widlund. Domain decomposition methods: algorithms and theory. Vol. 34. Springer, 2005.
- [21] AJ Cleary, RD Falgout, VE Henson, JE Jones, TA Manteuffel, SF McCormick, GN Miranda, JW Ruge, Robustness and Scalability of Algebraic Multigrid, SIAM J. Sci. Comput., 21, 2000, 1886–1908.
- [22] X. Cai, and M. Sarkis, A restricted additive Schwarz preconditioner for general sparse linear systems, SIAM Journal on Scientific Computing 21.2 (1999): 792-797.
- [23] AJ Cleary, RD Falgout, VE Henson, JE Jones, Coarse grid selection for parallel algebraic multigrid, in Proceedings of the fifth international symposium on solving irregularly structured problems in parallel, Springer-Verlag, New York, 1998.
- [24] A Sedighi, Y Deng, P Zhang, Fairness of Task Scheduling in High Performance Computing Environments, Scalable Computing: Practice and Experience, vol. 15, no. 3, pp. 273-285, 2014.
- [25] MT Jones, PE Plassman, A parallel graph coloring heuristic, SIAM Journal on Scientific Computing, 14(1993): 654-669.
- [26] M Luby, A simple parallel algorithm for the maximal independent set problem, SIAM Journal on Computing, 15(1986), 1036-1053.
- [27] JW Ruge and K Stüben, Algebraic multigrid (AMG), in: S.F. McCormick (Ed.), Multigrid Methods, Frontiers in Applied Mathematics, Vol. 5, SIAM, Philadelphia, 1986.
- [28] A Brandt, SF McCormick, J Ruge, Algebraic multigrid (AMG) for sparse matrix equations D.J. Evans (Ed.), Sparsity and its Applications, Cambridge University Press, Cambridge, 1984, 257–284.
- [29] RD Falgout, An Introduction to Algebraic Multigrid, Computing in Science and Engineering, Special Issue on Multigrid Computing, 8, 2006, 24–33.
- [30] K. Stüben, T. Clees, H. Klie, B. Lou, M.F. Wheeler, Algebraic multigrid methods (AMG) for the efficient solution of fully implicit formulations in reservoir simulation, SPE Reservoir Simulation Symposium. 2007.
- [31] K. Stüben, A review of algebraic multigrid, Journal of Computational and Applied Mathematics 128.1 (2001): 281-309.
- [32] Z. Chen, G. Huan, and Y. Ma. Computational methods for multiphase flows in porous media, Vol. 2. Siam, 2006.

- [33] R. D. Falgout, and U.M. Yang, HYPRE: A library of high performance preconditioners, Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2002. 632-641.
- [34] H. Liu, Dynamic Load Balancing on Adaptive Unstructured Meshes, 10th IEEE International Conference on High Performance Computing and Communications, 2008.
- [35] M.A. Christie, and M. J. Blunt, Tenth SPE comparative solution project: A comparison of upscaling techniques. SPE Reservoir Evaluation & Engineering 4.4 (2001): 308-317.
- [36] G. Karypis, K. Schloegel, and V. Kumar, Parallel static and dynamic multi-constraint graph partitioning, Concurrency and Computation: Practice and Experience, 2002, 14(3): 219-240.
- [37] S Lacroix, YV Vassilevski, MF Wheeler, Decoupling preconditioners in the implicit parallel accurate reservoir simulator (IPARS), Numerical linear algebra with applications, 8(8), 2001: 537-549.
- [38] B. Wang, S. Wu, Q. Li, X. Li, H. Li, C. Zhang, J. Xu, A Multilevel Preconditioner and Its Shared Memory Implementation for New Generation Reservoir Simulator, SPE-172988-MS, SPE Large Scale Computing and Big Data Challenges in Reservoir Simulation Conference and Exhibition, 15-17 September, Istanbul, Turkey, 2014.
- [39] K Wang, LB Zhang and Z Chen, Development of Discontinuous Galerkin Methods and a Parallel Simulator for Reservoir Simulation, SPE-176168-MS, SPE/IATMI Asia Pacific Oil & Gas Conference and Exhibition, 20-22 October, Nusa Dua, Bali, Indonesia, 2015.
- [40] H. Liu, K. Wang, Z. Chen, and K. Jordan, Efficient Multi-stage Preconditioners for Highly Heterogeneous Reservoir Simulations on Parallel Distributed Systems, SPE-173208-MS, SPE Reservoir Simulation Symposium held in Houston, Texas, USA, 23-25 February 2015.
- [41] H. D. Simon, Partitioning of unstructured problems for parallel processing, in: Proc. Conference on Parallel Methods on Large Scale Structural Analysis and Physics Applications, Pergamon Press, 1991.
- [42] A. Pothen, H. Simon, K. Liou, Partitioning sparse matrices with eigenvectors of graphs, SIAM J. Matrix Anal. 11 (3) (1990) 430-452.
- [43] T. Bui, C. Jones, A heuristic for reducing fill in sparse matrix factorization, in: Proc. 6th SIAM Conf. Parallel Processing for Scientific Computing, SIAM, 1993, pp. 445-452.
- [44] B. Hendrickson, R. Leland, A multilevel algorithm for partitioning graphs, in: Proc. Supercomputing 95, ACM, 1995.
- [45] G. Karypis, V. Kumar, A fast and high quality multilevel scheme for partitioning irregular graphs, Tech. Rep. CORR 95-035, University of Minnesota, Dept. Computer Science, Minneapolis, MN (June 1995).
- [46] P. M. Campbell and K. D. Devine and J. E. Flaherty and L. G. Gervasio and J. D. Teresco, Dynamic load balancing using space-filling curves, Technical Report CS-03-01, 2003.
- [47] J. G. Griffiths, Table-driven algorithms for generating space-filling curves. Computer-Aided Design 1985; 17(1): 37-41.

- [48] H. Liu, K. Wang, Z. Chen, K. Jordan, J. Luo, H. Deng, A Parallel Framework for Reservoir Simulators on Distributed-memory Supercomputers, SPE-176045-MS, SPE/IATMI Asia Pacific Oil & Gas Conference and Exhibition, Nusa Dua, Indonesia, 20-22 October, 2015.
- [49] C. Li, Y. Feng, Algorithm for analyzing n-dimensional Hilbert curve , vol. 3739. Springer Berlin/Heidelberg, 2005; 657–662.
- [50] N. Chen, N. Wang, B. Shi, A new algorithm for encoding and decoding the Hilbert order. *Software—Practice and Experience* 2007; 37(8): 897–908.
- [51] S. Kamata, R. O. Eason, Y. Bandou, A new algorithm for N-dimensional Hilbert scanning. *IEEE Trans on Image Processing* 1999; 8(7): 964–973.
- [52] X. Liu, G. F. Schrack, Encoding and decoding the Hilbert order. *Software—Practice and Experience* 1996; 26(12): 1335–1346.
- [53] X. Liu, G. F. Schrack, An algorithm for encoding and decoding the 3-D Hilbert order. *IEEE transactions on image processing* 1997; 6: 1333–1337.
- [54] H. Liu, Researches on Dynamic Load Balancing Algorithms and *hp* Adaptivity in 3-D Parallel Adaptive Finite Element Computations, PhD dissertation, Chinese Academy of Sciences, 2010.
- [55] A. J. Fisher, A new algorithm for generation hilbert curves. *Software: Practice and Experience* 1986; 16: 5–12.
- [56] A. R. Butz, Alternative algorithm for Hilbert’s space-filling curve. *IEEE Transactions on Computers* 1971; 20: 424–426.
- [57] L. M. Goldschlager, Short algorithms for space-filling curves. *Software—Practice and Experience* 1981; 11: 99–100.
- [58] I. H. Witten, B. Wyvill, On the generation and use of space-filling curves. *Software—Practice and Experience* 1983; 13: 519–525.
- [59] A. J. Cole, A note on space filling curves. *Software—Practice and Experience* 1983; 13: 1181–1189.
- [60] H. Sagan, *Space-Filling Curves*. Springer-Verlag; 1994.